ALCF Early Science Program

Argonne NATIONAL LABORATORY

ESP Kick-Off Workshop
Project Plan Presentation

Cosmic Structure Probes of the Dark Universe

PI: Salman Habib (LANL)

Presenter: Salman Habib

October 18–19, 2010

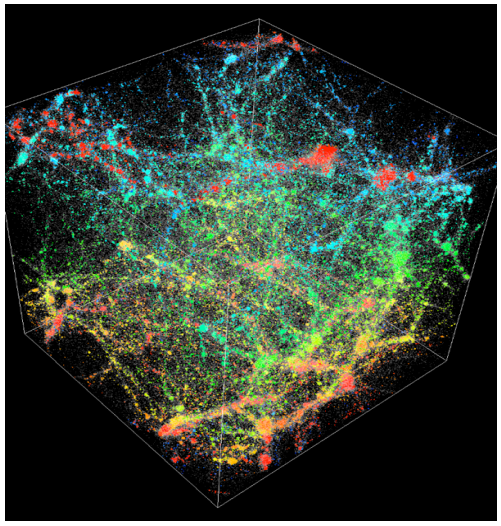

Project Overview

Structure Formation in the Universe as a Cosmological Probe

- **Contribute to characterization of dark energy and dark matter by predicting observational signatures for a variety of cosmological probes**
  - **First cosmological simulations resolving galaxy-scale mass concentrations at the scale of state-of-the-art sky surveys**
  - **Suite of simulations for 'cosmic calibration'**

- **Investigate primordial fluctuations by predicting the effects on cosmic structures at the current epoch**
  - **Systematic studies of running of the spectral index, primordial non-Guassianity**

Scientific Field: Cosmology

Modified HACC (Hardware-Accelerated Cosmology Code) Framework

Computational Approach, Numerical Methods

- **Hybrid particle/grid gravitational N-body code**
  - Solves the gravitational Vlasov-Poisson equation in an expanding universe
  - Large dynamic range in (i) force (1:1000000) and (ii) mass (1:10000)
  - Particle numbers of order 0.1-1 Trillion, FFTs in the 10000+ cubed class
  - Emphasis on accuracy (max errors of order 0.1-1%)
- **Inter-particle interaction broken into two, roughly equally numerically intensive parts**
  - Long/medium-range force
    - Grid-based, spectrally-filtered, Poisson solver (FFT-based)
  - Short-range force (two options)
    - Direct particle-particle interactions (with chaining mesh)
    - Tree algorithm
- **Symplectic time-stepping with adjustable time variable**
  - Global time-stepping with long-short sub-cycling
- **'On the Fly' analysis suite**
  - Fast halo and sub-halo finders
  - Statistical tools (correlation functions, power spectrum)
  - Weak gravitational lensing pipeline (under development)

Parallelism and Existing Implementation

- **Particle communication across nodes minimized using 'particle overloading' (memory cache):**
  - Node has particles extending significantly beyond spatial subdomain of node (rough analog of 'ghost zones')
    - Copies of these "passive" particles held by other nodes
    - Single particle only "alive" in one node
  - Many steps require no particle communication
  - Refresh overloading zones as needed, only nearest-neighbor communication needed
- **Digital filtering/differencing in the spectral domain**
  - Simple communication templates
- **MPI plus Cell BE (Roadrunner)**
  - 3D domain decomposition
- **MPI plus OpenCL (GPU clusters)**
- **Current Performance/Scalability**
  - Weak scaling across entire Roadrunner
  - Large fraction of peak performance

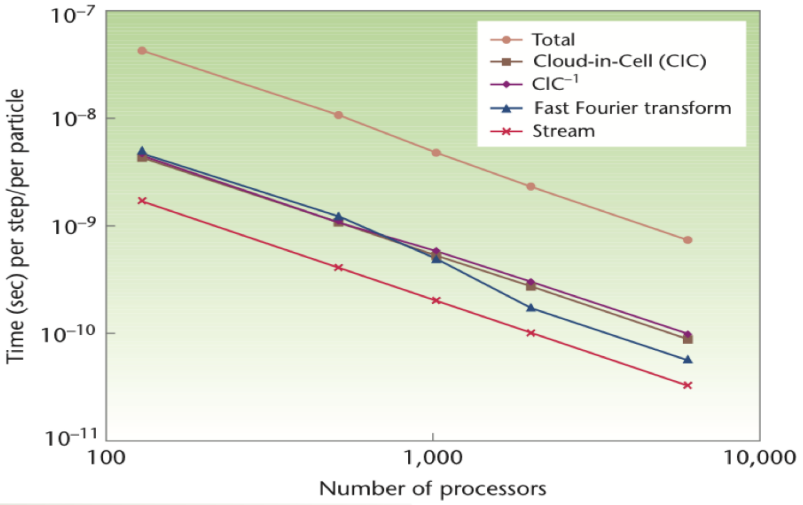

Library and Tool Dependencies

- **'Lightweight' code with minimal dependencies (except for 'standards' such as MPI, FFTW, --)**
- **Libraries**
  - **FFTW for serial FFTs (ESSL alternatives would be useful)**
  - **HDF (I/O) would be useful, although not used currently (n-n read/writes)**
- **Tools**
  - **ParaView for parallel visualization and visualization-aided analysis (co-developed with LANL Paraview team)**

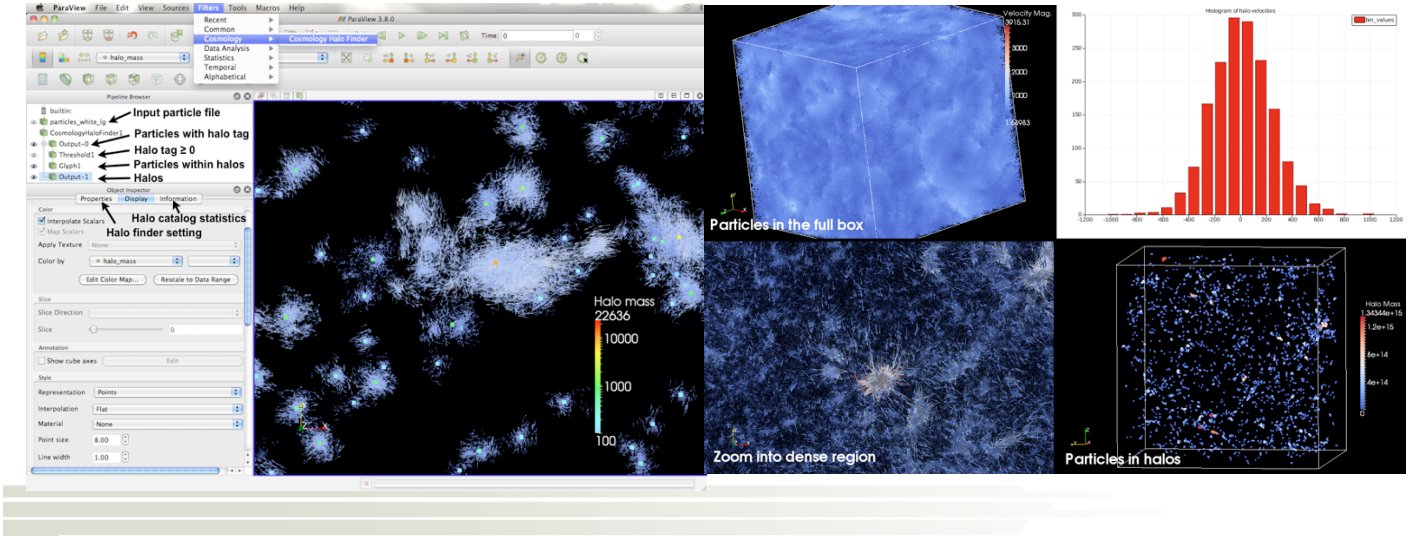

Anticipated Modifications for Blue Gene/Q

- **No change needed for grid layer of code**
  - Performance of MPI with parallel FFT should be adequate on Mira
- **Node-level modifications needed**
  - Current code optimized for Cell-based (CPU-accelerator memory-balanced) nodes on Roadrunner and for GPU clusters with memory-imbalanced nodes
  - Rewrite to exploit threads and instructions on BG/Q core, test both particle-particle and tree-based algorithms for the short-range force-solver
  - Optimize real-time analysis framework for Mira
- **Performance and scaling needed to run proposed problem on Mira:**
  - State-of-the-art sky survey volume: Order of 1-100 billions of Mpc^3
    - Of order 0.1-1 trillion grid cells
    - Of order 0.1-1 trillion particles
  - Wall-clock of order 2 minutes per long (sub-cycled) time-step
  - Keep real-time analysis 'slots' down to of order 10 minutes/slot

Plan for Next 6 Months Effort

- **Help find and hire a project postdoc**
- **Port to Blue Gene/P**
  - Complete port of long/medium-range-only version (mostly working)
  - Simultaneously start work on local force computation
- **Method selection for hydro capability (fixed-mesh)**
  - At smaller length scales, hydrodynamic and feedback effects from baryons significant, need this capability down the road
  - Particle-based methods likely to scale better on next-generation architectures, investigate modern hydro-PIC as alternative to SPH
- **Benchmark long-range code performance**
  - Use projections to estimate performance on Blue Gene/Q (mostly controlled by performance of parallel 3-D FFT)